

X-ray Spectra from MHD Simulations of Accreting Black Holes

Jeremy D. Schnittman

*NASA Goddard Space Flight Center
Greenbelt, MD 20771*

`jeremy.schnittman@nasa.gov`

and

Julian H. Krolik

*Department of Physics and Astronomy,
Johns Hopkins University
Baltimore, MD 21218*

`jhk@pha.jhu.edu`

and

Scott C. Noble

*Center for Computational Relativity and Gravitation,
Rochester Institute of Technology
Rochester, NY 14623*

`scn@astro.rit.edu`

ABSTRACT

We present the results of a new global radiation transport code coupled to a general relativistic magneto-hydrodynamic simulation of an accreting, non-rotating black hole. For the first time, we are able to explain from first principles in a self-consistent way the X-ray spectra observed from stellar-mass black holes, including a thermal peak, Compton reflection hump, power-law tail, and broad iron line. Varying only the mass accretion rate, we are able to reproduce the low/hard, steep power-law, and thermal-dominant states seen in most galactic black hole sources. The temperature in the corona is $T_e \sim 10$ keV in a boundary layer near the disk and rises smoothly to $T_e \gtrsim 100$ keV in low-density regions far above the disk. Even as the disk's reflection edge varies from the horizon out to

$\approx 6M$ as the accretion rate decreases, we find that the shape of the Fe $K\alpha$ line is remarkably constant. This is because photons emitted from the plunging region are strongly beamed into the horizon and never reach the observer. We have also carried out a basic timing analysis of the spectra and find that the fractional variability increases with photon energy and viewer inclination angle, consistent with the coronal hot spot model for X-ray fluctuations.

Subject headings: black hole physics – accretion disks – X-rays:binaries

1. INTRODUCTION

Since the initial discovery of the magneto-rotational instability (MRI) by Balbus & Hawley (1991) over two decades ago, tremendous progress has been made in simulating astrophysical accretion disks. Large-scale magneto-hydrodynamic (MHD) simulations have steadily improved their resolution and physical accuracy, leading to greater understanding of the fundamental physics of accretion. Yet despite all this success, we are not much closer to reproducing X-ray observations of accreting black holes than we were after the original papers on the structure of their surrounding disks (Shakura & Sunyaev 1973; Novikov & Thorne 1973). From basic conservation laws and a guessed boundary condition, it is relatively easy to explain the thermal spectrum seen in some black holes with a multi-color disk model (Mitsuda et al. 1984). However, it was realized very early on that a great deal of the power of stellar-mass black holes and active galactic nuclei (AGN) is in the form of high-energy X-rays well above the thermal peak (Oda et al. 1971; Elvis et al. 1978).

Although it is now widely accepted that this hard flux comes from the inverse Compton scattering of seed photons from the disk through a hot corona (Liang & Price 1978; Haardt & Maraschi 1993), we still know little or nothing about the origin or detailed properties of this corona. Almost all previous work has been largely phenomenological and guided by energy conservation, using parameterized models to divide the total dissipation between the disk and corona in an attempt to best fit the data (Svensson & Zdziarski 1994; Done & Kubota 2006). Numerous papers have shown for galactic black holes and AGN that the hard spectrum requires the coronal heating to be spatially localized and inhomogeneous (Haardt et al. 1994; Stern et al. 1995; Zdziarski et al. 1996; Poutanen et al. 1997), but there are many ways in which this can happen.

Global MHD simulations are ideal for understanding the dynamics of the corona from first principles, but until now we have lacked the tools necessary for closing the loop and comparing the results of the simulations directly with the observations they are meant to

explain. In this paper we attempt to bridge the gap between theory and observation by applying the radiative transfer code *Pandurata* (Schnittman & Krolik 2012) as a “post-processor” to simulation data generated by the general relativistic MHD code *Harm3d* (Noble et al. 2009). In this way, *we are able to explain from first principles the origin of the hard X-ray flux in a self-consistent way.*

Whereas analytic disk models rely on dimensional analysis to describe the scaling of shear stress with pressure, direct calculation of the nonlinear development of the MRI allows us to compute quantitatively the rate of angular momentum transfer through the actual magnetic stresses. Thus, we can dispense with the greatest uncertainties of the traditional accretion model: the dependence of stress on local physical conditions; the spatial distribution of dissipation; and the inner boundary condition, which can now be moved *inside* the horizon, and thus made physically irrelevant. Local energy dissipation can be easily monitored with the flux-conservative code *Harm3d*, which contains an heuristic cooling function designed to generate thin disks (Noble et al. 2009). While Shakura & Sunyaev (1973) relied on steady-state, azimuthally- and vertically-averaged quantities, the MHD simulations provide dynamic, three-dimensional information about the fluid density, 4-velocity, magnetic pressure, gas pressure, and cooling at every point throughout the computational domain.

The first step on the path from simulation to observation is to convert the code variables to physical units by specifying the black hole mass and accretion rate. We can then distinguish between the optically thick disk body and the optically thin corona. At this point, we make an assumption similar to one made by Shakura & Sunyaev (1973) and Novikov & Thorne (1973), that energy dissipated in the disk is emitted as thermal radiation from the disk surface. After demonstrating that other potentially relevant emission mechanisms (e.g., thermal bremsstrahlung, synchrotron) are negligible, we balance local energy dissipation and inverse Compton up-scattering of disk seed photons in order to find the equilibrium electron temperature and radiation intensity as functions of position throughout the corona.

As described in a companion paper (Schnittman & Krolik 2012), *Pandurata* is a fully relativistic Monte Carlo radiation transport code that integrates photon trajectories from the disk surface, accounts for scattering through the hot corona, and transports them to their ultimate destination, either a distant observer or the black hole horizon. Although *Pandurata* includes polarization effects in all its scattering calculations, the results in this paper focus on spectral features alone. For stellar-mass black holes and Eddington-scaled luminosities $\dot{m} = L/L_{\text{Edd}} \approx 0.01 - 1.0$, we find the simulated X-ray continuum spectra comprise a thermal peak around 1 keV, a Compton reflection hump from 30–100 keV, and a power-law tail extending up to $\gtrsim 1$ MeV. By varying the mass accretion rate, we can reproduce the three main accretion states described in Remillard & McClintock (2006):

hard, thermal, and steep power law. Although the results in this paper are based entirely on a single simulation, whose structure most closely matches the classical predictions for a disk with $\dot{m} \approx 0.1 - 0.3$, there is qualitative agreement with observations spanning the entire range of $\dot{m} = 0.01 - 1.0$.

We have also included a simple model for fluorescent line production and can reproduce the Fe $K\alpha$ features seen in many galactic black holes and AGN. We find effective widths of $\approx 200 - 350$ eV, consistent with observations (Miller et al. 2004, 2006a; Walton et al. 2012). For very low accretion rates, the vertically-integrated optical depth falls below unity in the inner regions, so the disk is effectively truncated around $r \approx 4 - 6M$. Interestingly, the iron line profile appears independent of the location of the reflection edge, as long as it is inside the inner-most stable circular orbit (ISCO). Lastly, we include some rudimentary variability analysis, finding results consistent with a large body of observations: the low-hard state is more variable than the disk-dominated state, and in all states the fractional RMS amplitude increases with photon energy (Remillard & McClintock 2006). On short time scales, the amplitude of fluctuations increases with observer inclination angle, consistent with the coronal hot spot model of X-ray variability.

2. SIMULATION DATA


2.1. Description of Harm3d

The data we analyze for this paper are drawn from the highest resolution simulation reported in Noble et al. (2010) and Noble et al. (2011), designated “ThinHR” in those papers. Harm3d, the code used to generate the data, is an intrinsically conservative 3D MHD code in full general relativity; this particular simulation was computed in a Schwarzschild spacetime. Because it uses a coordinate system based on Kerr-Schild, Harm3d is able to place the inner boundary of the computation volume inside the black hole’s event horizon, thus obviating the need for any guessed inner boundary conditions. The stress-energy conservation equation is modified to include a local cooling function; that is, we write $\nabla_\nu T_\mu^\nu = -\mathcal{L}u_\mu$, where T_μ^ν is the stress-energy tensor, u_μ is the specific 4-momentum, and $\mathcal{L}$ is non-zero only for gravitationally-bound gas, and only when the local temperature is greater than a target temperature T_* . When the temperature exceeds that threshold, the excess heat is radiated away on an orbital timescale. The target temperature T_* is chosen so as to keep the disk’s aspect ratio H/r close to a single pre-set value at all radii. In dimensionless code units, $T_* \equiv (\pi/2)(R_z/r)(H/r)^2$, where R_z describes the correction to the vertical gravity due to relativistic effects (Noble et al. 2010). For ThinHR simulation, the target scale height was $H/r = 0.06$.

We took special pains to ensure the numerical quality of these simulations. Every $20M$ in time (we set $G = c = 1$, so time has units of $(M/M_\odot) \cdot 4.9 \times 10^{-6}$ s, and distance has units of $(M/M_\odot) \cdot 1.5 \times 10^5$ cm), we measured the number of cells across the fastest growing MRI wavelength in both the vertical and the azimuthal directions (λ_z and λ_ϕ). The minimum number to achieve the correct linear growth rate for vertical modes is 6 cells per λ_z (Sano et al. 2004); to describe nonlinear behavior, at least 20 cells per λ_ϕ and at least 10 per λ_z are necessary (Hawley et al. 2011). The mass- and time-weighted values in ThinHR were 25 (vertical) and 18 (azimuthal). As discussed in Hawley et al. (2011), by this and several other measures, ThinHR is the best-resolved global accretion simulation in the literature. By examining the time-dependent hydrodynamic and radiative properties of the fluid at several fiducial radii, Noble et al. (2010) determined that the final $5000M$ of the ThinHR simulation met the relevant criteria for inflow equilibrium in the inner disk. We therefore restrict our analysis of the simulation data to that period throughout this paper.

In studying simulations intended to represent statistically steady accretion, it is important to recognize that when there is only a finite amount of mass on the grid, some of it must move out in order to absorb the angular momentum removed from accreted material. Consequently, the radial range over which the disk can be said to be in inflow equilibrium is limited. For the simulations under consideration here, that range was typically $r \lesssim 20M$. In Noble et al. (2011), the disk beyond this radius was simply replaced with a standard relativistic Novikov-Thorne (N-T) thin disk (Novikov & Thorne 1973). That paper focused exclusively on thermal radiation, so a thin disk was an appropriate extrapolation of the simulation data beyond $20M$. Here, we are primarily interested in the coronal properties of the accretion flow, for which there are no simple analytic solutions. Therefore, we include the entire body of simulation data out to $\sim 60M$, beyond which the surface density of the gas begins to decrease rapidly, and the accretion geometry no longer resembles a disk, but rather a torus. Despite the fact that the disk is not strictly in inflow equilibrium outside of $\sim 20M$, the dissipation profile still roughly follows that expected for a classical disk. This can be seen in Figure 1, where we plot the radial shell-integrated dissipation profile for the Harm3d data, along with that given by N-T for the same accretion rate. Unlike Noble et al. (2011), here we make no attempt to normalize the dissipation profile by the radial mass accretion rate, which explains the somewhat larger deviation of Harm3d from N-T outside $\sim 10M$ shown in this figure than in Figure 2 of Noble et al. (2011).

Fig. 1.— Luminosity profile $dL/d(\log r)$, integrated over θ and ϕ , and averaging over time. The solid curve is the Harm3d data, and the dashed curve is the Novikov-Thorne prediction. For both cases, the Eddington-normalized accretion rate is $\dot{m} = 0.1$.



2.2. Conversion from Code to Physical Units

When comparing the Harm3d predictions with real physical systems, the first step is always to convert the fluid variables from dimensionless code units to physical cgs units. This conversion requires specifying the black hole mass M , which sets the natural length and time scales, and the accretion rate $\dot{M}$, which determines the scale for the gas density, cooling rate, and magnetic pressure. One technique for enacting this conversion is described in the appendix of Noble et al. (2011), where the actual ray-tracing calculation is done in dimensionless units, and only the final observed spectrum is converted to physical units. That approach works best for optically thin systems where a single emission mechanism is used throughout the accretion flow and the photons do not interact with the matter.

In this work, however, we are interested in very different radiation processes in the disk and the corona. Therefore, before we even begin the ray-tracing calculation, a photosphere must be defined to distinguish between the cool, dense disk and the relatively hot, diffuse corona. Since the location of this photosphere is a function of the gas density, it will be different for different values of the accretion rate. Following Schnittman et al. (2006) and Noble et al. (2009), the physical density is given by

$$\rho_{\text{cgs}} = \rho_{\text{code}} \frac{4\pi c^2}{\kappa GM} \frac{\dot{m}/\eta}{\dot{M}_{\text{code}}}, \quad (1)$$

where $\kappa = 0.4 \text{ cm}^2/\text{g}$ is the electron scattering opacity, and $\dot{m}$ is the Eddington-scaled accretion rate assuming a radiative efficiency $\eta = 0.06$. For ThinHR, the mean accretion rate in code units was $\dot{M}_{\text{code}} = 3 \times 10^{-4}$. Regardless of the exact value for η used in the conversion from code units to cgs, the ray-tracing procedure itself results in an independent value for the radiative efficiency, which is listed in Table 1.

Table 1: For a range of mass accretion rates: the bolometric radiative efficiency η , the time-averaged fraction of total luminosity in the corona, the radius of the reflection edge R_{refl} , the disk-corona transition radius R_{trans} , and the height H_{phot} of the scattering photosphere. The dependence of η on $\dot{m}$ is in part an artifact of our model, as explained in the text. Note also that emission outside $R = 60M$, ignored here, adds an additional 0.012 to the radiative efficiency.

$\dot{m}$	η	$L_{\text{cor}}/L_{\text{tot}}$	R_{refl}/M	R_{trans}/M	H_{phot}/r
0.01	0.056	0.40	6.1	8.8	0.11
0.03	0.052	0.29	4.4	7.4	0.19
0.1	0.051	0.19	2.1	6.4	0.31
0.3	0.048	0.13	2.0	5.7	0.43
1.0	0.042	0.09	2.0	5.1	0.55


Once the physical density is specified, the location of the photosphere at each point in the disk at any particular time is calculated by integrating the optical depth $d\tau = \kappa \rho(r, \theta, \phi) r d\theta$ at constant (r, ϕ) from the poles at $\theta = 0, \pi$ down towards the disk. The photosphere is then defined as the surface where the integrated optical depth reaches unity. For the top and bottom of the disk, the photosphere surfaces can be written as $\Theta_{\text{top}}(r, \phi)$ and $\Theta_{\text{bot}}(r, \phi)$ as in Schnittman & Krolik (2012):

$$\int_{\theta=0}^{\theta=\Theta_{\text{top}}} d\tau = \int_{\theta=\Theta_{\text{bot}}}^{\theta=\pi} d\tau = 1. \quad (2)$$

Figure 2 shows a snapshot of the gas density in the $(r, z = r \cos \theta; \phi = 0)$ plane for fiducial values of the black hole mass $M = 10M_{\odot}$ and accretion rate $\dot{m} = 0.1$. The solid contour lines show surfaces of constant optical depth. Note that while the density-weighted scale height of the disk H/r is only ≈ 0.06 , the photosphere is located at several scale heights above the midplane, with $H_{\text{phot}}/r \approx 0.3$ for this choice of accretion rate. This is to be expected; in stratified shearing box simulations with careful treatment of thermodynamics and radiation transfer, the scattering photosphere often lies 3–4 scale heights from the plane (Hirose et al. 2009a).

For $\dot{m} = 0.1$, the total optical depth of the disk ranges from order unity in the plunging region up to $\tau \approx 100 - 200$ in the disk body at $r > 10M$. Where the total optical depth is less than 2, we say that there is no disk, only corona (i.e., no solution exists for eqn. 2). We denote the radius of this transition by R_{ref} ; in the language of Krolik & Hawley (2002), this is the radius of the “reflection edge.”

Fig. 2.— Fluid density profile for a slice of Harm3d data in the (r, z) plane at simulation time $t = 12500M$. Contours show surfaces of constant optical depth with $\tau \approx 0.01, 0.1, 1.0$. Fiducial values for the black hole mass $M = 10M_{\odot}$ and accretion rate $\dot{m} = 0.1$ were used.




Just as the gas density must be converted from code units to physical units, so do the magnetic field and local cooling rate. With dimensional analysis, determining these conversion factors is trivial. In cgs units, the magnetic energy density is given by $U_B = B^2/(8\pi)$, so the conversion factor is simply

$$\frac{B_{\text{cgs}}^2}{B_{\text{code}}^2} = c^2 \frac{\rho_{\text{cgs}}}{\rho_{\text{code}}}. \quad (3)$$

The local cooling rate $\mathcal{L}$ has units of energy density per time, so its conversion factor is given by


$$\frac{\mathcal{L}_{\text{cgs}}}{\mathcal{L}_{\text{code}}} = c^2 \frac{\rho_{\text{cgs}}}{\rho_{\text{code}}} \frac{t_{\text{code}}}{t_{\text{cgs}}} = \frac{c^5}{GM} \frac{\rho_{\text{cgs}}}{\rho_{\text{code}}} \quad (4)$$

Fig. 3.— Magnetic energy density profile for a slice of Harm3d data in the (r, z) plane corresponding to the same conditions as in Figure 2.



In Figures 3 and 4 we show the magnetic energy density and local cooling function, respectively. The Harm3d data correspond to the same time and the same slice in the (r, z) plane as shown in Figure 2, for $M = 10M_{\odot}$ and $\dot{m} = 0.1$. Comparing the gas density and magnetic pressure, we see that both are concentrated in the disk, but the magnetic scale height is significantly greater than the gas scale height. This contrast naturally leads to a corona that is dominated by magnetic pressure, as seen in most shearing box and global MHD simulations. From equation (1), we see that the physical density scales inversely with black hole mass, but the physical length scale is proportional to M , so the location

Fig. 4.— Local cooling rate for a slice of `Harm3d` data in the (r, z) plane corresponding to the same conditions as in Figure 2. Black regions contribute zero emission.




of the photosphere—and thus the relative fraction of power from the disk and corona—is independent of M .

The cooling profile appears to closely follow the magnetic field, consistent with earlier models that use magnetic stress as a proxy for heat dissipation (Beckwith et al. 2008), as well as stratified shearing box simulations in which the actual dissipation rate is computed (Hirose et al. 2006). As described above, `Harm3d` uses a local cooling function $\mathcal{L}$ to keep the disk relatively thin. This cooling can also be thought of as the local dissipation of heat, so we will often identify $\mathcal{L}$ as the emissivity of the gas. Because at any given time some of the fluid elements are actually below their target temperatures, the contours of $\mathcal{L}$ show numerous isolated patches with no emission (black in Fig. 4). In fact, the vast majority of the coronal emission comes from a relatively small volume of space. Figure 5 shows the cumulative fraction of the coronal volume responsible for the cumulative fraction of the

coronal luminosity: 50% of the corona volume generates only 1% of the luminosity, while 10% of the corona generates 90% of the luminosity, consistent with many earlier models that assume a highly inhomogeneous heating profile in the corona (Haardt et al. 1994; Stern et al. 1995; Zdziarski et al. 1996; Poutanen et al. 1997).


Fig. 5.— Fraction of the coronal volume that generates a given fraction of the total coronal luminosity, for $\dot{m} = 0.1$. 10% of the corona is responsible for 90% of the emissivity.



As can be seen from Figure 4, for $\dot{m} = 0.1$, the majority of the total emission comes from within the disk body, with a sizable contribution from the corona in the innermost regions. Outside a transition radius $R_{\text{trans}} \approx 6M$, the shell-integrated luminosity is dominated by the disk, while inside of this radius the corona dominates. By increasing $\dot{m}$ for the same simulation, the density scale increases [see eqn. (1)], encompassing a greater fraction of the total luminosity within the optically thick disk. Conversely, for small values of $\dot{m}$, the disk shrinks and the corona becomes more dominant. Table 1 shows how the relative contributions of the disk and corona change with $\dot{m}$, as well as the locations of R_{ref} and R_{trans} . The corona dominates throughout the plunging region for all sub-Eddington values of $\dot{m}$. This can also be seen in Figure 6, where we have plotted the fraction of total dissipation (angle- and time-averaged) occurring in the corona at each radius for a range of $\dot{m}$. This fraction can be directly compared to the f parameter used in the coupled disk-corona model of Done & Kubota (2006), where the total dissipation at each radius is divided between disk and corona. Thus it should come as no surprise that our resulting spectra (see below in Sec. 4)

are qualitatively similar to those that they predicted for comparable values of f .

Fig. 6.— Fraction of total dissipation in the corona as a function of radius, for a range of accretion rates $\dot{m}$.



With increasing $\dot{m}$, the photosphere height naturally increases, making the disk more like a bowl or inverted cone. This shape increases the probability of scattering off the disk surface (possibly losing energy in the process) and being captured by the black hole. Thus, the radiative efficiency decreases steadily with larger $\dot{m}$. This effect may be interpreted as the beginning of “super-Eddington photon trapping”, but likely underestimates the true efficiency because our disk model forces $H \propto r$, rather than the more physically realistic $H \simeq \text{const}$. Note that even the thinnest disks modeled here have lower η than quoted in Noble et al. (2011) because we ignore any emission outside $60M$. For standard N-T disks, the emission beyond this radius adds an additional $\delta\eta = 0.012$.

2.3. Global Radiation Field

Given the density, internal energy, and an equation of state, it is possible to derive a gas temperature directly from the simulation data. However, because `Harm3d` does not include

any radiation pressure in its equation of state, this inferred temperature would not be very meaningful. Inside the optically thick disk, the actual gas temperature would be lower in proportion to the degree of radiation pressure support. In the corona, this method would at best yield an ion temperature rather than an electron temperature, and thus would be more or less irrelevant for radiation processes (see the discussion at the end of Sec. 3 to a more detailed analysis of this issue). Therefore, the electron temperature must be derived from other simulation data through physical reasoning more directly tied to their heating and cooling.

Inside the disk (i.e., between the two electron scattering photospheres), we assume all the emitted radiation is able to thermalize, and all the heat generated within the disk is radiated from the same (r, ϕ) where it is made. These two assumptions allow us to define an effective temperature at each point on the disk surface:

$$\sigma T_{\text{eff}}^4(r, \phi) = \frac{1}{2} \int_{\Theta_{\text{top}}}^{\Theta_{\text{bot}}} \mathcal{L}(r, \theta, \phi) d\ell, \quad (5)$$

where σ is the Stefan-Boltzmann constant and the factor of $1/2$ is due to the fact that half the radiation is emitted from the top and half from the bottom of the disk.

In the regions of the disk where most of the flux is generated, electron scattering opacity is always much greater than the opacity due to other processes such as free-free absorption (Shakura & Sunyaev 1973). We therefore assume the mean intensity in the fluid frame at the photosphere has the spectrum of a diluted black-body:

$$J_\nu = f_{\text{hard}}^{-4} B_\nu(f_{\text{hard}} T_{\text{eff}}), \quad (6)$$

where $B_\nu(T)$ is the Planckian black-body function, and the hardening parameter f_{hard} is taken to be 1.8 (Shimura & Takahara 1995). As described in Schnittman & Krolik (2012), the disk intensity also has an angular dependence (limb-darkening) and polarization (Chandrasekhar 1960).

In the corona, the picture is not so simple. Unlike in the disk, in the corona the radiation is not expected to be thermalized. We have considered the contributions from a number of different radiation mechanisms, including bremsstrahlung, cyclo/synchrotron, and inverse Compton (IC), but, at least for the stellar-mass black holes of interest here, we find the coronal power is completely dominated by inverse Compton.

Bremsstrahlung and synchrotron emission are both fundamentally local processes, and depend only on the local electron density, temperature, and magnetic field. Since the density, magnetic field, and net emissivity are given by the Harm3d data, to solve for the temperature

we could simply invert the following equation at each point in the corona if they were the only cooling agents:

$$\mathcal{L} = P_{\text{brem}}(\rho, T_e) + P_{\text{synch}}(\rho, T_e, B), \quad (7)$$

where P_{brem} and P_{synch} are the local bremsstrahlung and synchrotron power density (in erg/s/cm³), respectively. But this approach is incomplete because the corresponding absorptive opacity

$$\alpha_\nu = \frac{j_\nu}{B_\nu(T_e)} \quad (8)$$

can also be important. Here α_ν and j_ν are the absorption and emission coefficients respectively for either bremsstrahlung or synchrotron. For typical coronal conditions of $T_e \sim 10\text{--}1000$ keV, $n_e \sim 10^{16}\text{--}10^{18}$ cm⁻³, and $B \sim 10^6\text{--}10^7$ G, we find that free-free emission and absorption are both negligible, while synchrotron emission can actually contribute a significant fraction of the total cooling function. However, the typical cyclotron frequency for these parameters lies in the infrared, where self-absorption is strong. Since the corona is optically thick to synchrotron radiation, it does not end up contributing significantly to the total cooling: every photon that is emitted is almost instantly re-absorbed. We are thus left with IC as the dominant emission process in the corona.

Unlike bremsstrahlung or synchrotron, IC is a fundamentally *non-local* process because it requires a population of seed photons to be up-scattered by the hot electrons. Moreover, the IC seeds can come from distant parts of the accretion disk. Local treatments are therefore insufficient.

For a mono-energetic population of electrons with isotropic velocity v , the IC power is (Rybicki & Lightman 2004)

$$P_{\text{IC}} = \frac{4}{3} \sigma_T c \gamma^2 \beta^2 n_e U_{\text{ph}}. \quad (9)$$

Here σ_T is the Thomson cross section, $\beta = v/c$, $\gamma = (1 - \beta^2)^{-1/2}$, n_e is the electron density, and U_{ph} is the energy density of the local photon distribution. This local photon density is not known *a priori* from the simulation data, so it must be solved for using radiation transport. Because the corona has an optical depth of order unity (i.e., neither optically thin nor optically thick approximations can be used) and its geometry is complex, we use Monte Carlo ray-tracing including scattering to model the transport in a global manner (Schnittman & Krolik 2012).

With the assumption that the photon diffusion time is short compared to the time required for any dynamical or thermal changes, the problem can be thought of as a boundary value problem: given the B-field, density, fluid 4-velocity, and cooling rate at every point in the corona, along with the thermal seed photon distribution at the photosphere surface, we

need to solve for the electron temperature $T_e(r, \theta, \phi)$ and photon energy density $U_{\text{ph}}(r, \theta, \phi)$ at every point in the corona. To do so, we employ an iterative technique as follows:

- Initially estimate the local value of the radiation density in terms of the thermal contribution at the surface of the disk: $U_{\text{ph}}(r, \theta, \phi) = c\sigma T_{\text{disk}}^4(r, \phi)$.
- Solve equation (9) with $P_{\text{IC}} = \mathcal{L}$ to get $\gamma(r, \theta, \phi)$ throughout the corona. Derive the electron temperature at each point from the relation $T_e = \frac{2}{3} \frac{m_e c^2}{k_B} (\gamma - 1)$.
- Carry out a complete Monte Carlo ray-tracing calculation with **Pandurata**, using thermal seed photons from the disk photosphere propagating through the corona via Compton scattering.
- For each volume element in the corona, determine the total amount of IC power generated in that zone by comparing the ingoing and outgoing energy of every photon packet that scatters within that zone.
- Compare the coronal power from **Pandurata** with the cooling function $\mathcal{L}(r, \theta, \phi)$ given by **Harm3d**. Where the coronal power from the ray-tracing calculation exceeds the cooling rate in the simulation, the initial guess for U_{ph} was too low, giving a T_e that is too high, and vice versa.
- Revise the coronal temperature estimates up or down accordingly, and repeat the full ray-tracing calculation, getting a new 3D map of the cooling function, which is again compared with the target values from the simulation data.
- Repeat until the global solution for T_e , U_{ph} , and $\mathcal{L}$ is self-consistent throughout the corona.

Because the Monte Carlo technique is inherently noisy, the **Pandurata** calculation and the **Harm3d** target cooling rate for any individual fluid element are unlikely to agree very well. We typically use $\sim 10^7$ – 10^8 photon packets for a single snapshot, while the simulation volume contains roughly 10^6 cells. Consequently, for the majority of the corona volume, where $\tau < 0.1$, a given cell might see an average of only *one* photon packet. It should therefore not be surprising that point-to-point Poisson fluctuations are quite large. Furthermore, as can be seen in Figure 4, even the target cooling function is highly non-uniform, characterized by large-amplitude fluctuations on a small spatial scale.

For this reason, before comparing the results of the ray-tracing calculation with the target **Harm3d** emissivity, we apply a smoothing kernel to both data sets to remove the fluctuations described above. This smoothing is not only numerically helpful, but is also

strongly motivated from a physical point of view. For any radiation transport problem in a roughly steady state, the photon energy and momentum density cannot change significantly over length scales much shorter than the mean free path. Thus, when testing for convergence of U_{ph} , it is eminently reasonable to smooth the cooling function $\mathcal{L}$ over the characteristic scattering length.

In fact, by smoothing over an even greater volume, we can significantly improve the efficiency of our iterative solver. For example, if the smoothing length is comparable to the coronal scale height, then instead of trying to sample $\mathcal{L}$ in 10^6 fluid elements, we are effectively only probing 10–100, and thus can use many fewer photon packets. After converging at low resolution, we repeat the calculation with more photons and a shorter smoothing length until *Pandurata* and *Harm3d* agree to high accuracy everywhere.

One way to see this agreement at a quantitative level is to compare the radial distribution of coronal emission $dL/d(\log r)$ as derived from the two codes for a single snapshot, shown in Figure 7. After only 3 levels of iteration, we are clearly able to resolve coronal hot spots as small as $\Delta r/r \sim 0.2$. By plotting $dL/d(\cos \theta)$, we see that the vertical profile of the corona is also well-matched (see Fig. 8).


3. CORONAL TEMPERATURE

The global solution for the electron temperature is shown in Figure 9. The temperature within each vertical slice of the disk body is constant, and of the order 0.2–1 keV for these parameters. The corona is, of course, much hotter, with T_e ranging from ~ 10 –100 keV for τ between 0.01 and 1. By comparing Figures 2 and 4, we see that at high latitudes, the electron density falls off faster with increasing altitude from the disk than the dissipation, leading to higher coronal temperatures as more power must be released by a smaller quantity of gas. The temperature map also shows large fluctuations over small spatial scales, yet not quite as large as those seen in $\mathcal{L}$. This is because the regions of high dissipation are correlated with regions of high density, which has the effect of smoothing out the temperature gradients (see eqn. (9)).

By changing the $\dot{m}$ used in equations (1) and (4), we can investigate the coronal properties of different accretion states. Since $\mathcal{L}$, n_e , and U_{ph} all scale linearly with $\dot{m}$, from equation (9) one can see that the term $\gamma^2 \beta^2$ should scale like $\dot{m}^{-1}$. At low electron temperatures we have

$$\gamma^2 \beta^2 \approx \frac{v^2}{c^2} \approx 3 \frac{k_B T_e}{m_e c^2}, \quad (10)$$

Fig. 7.— Instantaneous luminosity profile $dL/d(\log r)$, integrated over θ and ϕ , considering only coronal emission. The black curve is the **Harm3d** data, and the red curve is the **Pandurata** ray-traced reconstruction. The luminosity is $0.1L_{\text{Edd}}$. The Monte Carlo data used $\sim 5 \times 10^7$ photon packets for this snapshot.




while in the relativistic regime,

$$\gamma^2 \beta^2 \approx \gamma^2 \approx \frac{9}{4} \left(\frac{k_B T_e}{m_e c^2} \right)^2, \quad (11)$$

recovering the well-known scaling of IC power with temperature (Rybicki & Lightman 2004). At low coronal temperatures, $T_e \sim \dot{m}^{-1}$, while at high temperatures, $T_e \sim \dot{m}^{-1/2}$, *independent of the black hole mass*.

This simple scaling of temperature with accretion rate is somewhat complicated by the fact that at different accretion rates, the relative fraction of power from the corona and from the disk also changes (see Table 1), but we do see a clear qualitative trend that is consistent with decades of observations: low-luminosity states are characterized by hard X-ray flux from a hot corona, while high-luminosity states lead to a much cooler corona and softer spectrum. In Figure 10 we plot the time-averaged coronal temperature as a function of radius for a range of different accretion rates. In the top panel the mean temperature is


Fig. 8.— Coronal luminosity profile as in Fig. 7, but for $dL/d(\cos\theta)$.



calculated by integrating over θ and ϕ and weighting by the local cooling rate $\mathcal{L}$, while in the bottom panel the temperature is weighted by the electron density n_e . The $\mathcal{L}$ -weighting is more closely related to the emergent spectrum and naturally probes the upper corona, while the n_e -weighting speaks to conditions in the majority of the coronal mass and is sensitive to the conditions near the disk. In either case, the trend with $\dot{m}$ is clear, and we also see that in the bulk of the corona, especially outside the ISCO, the temperature changes very little with radius.

The time-averaged radial and vertical temperature profiles of the corona can be seen in greater detail in Figure 11 for $\dot{m} = 0.1$. At six different values of r , we plot the temperature as a function of optical depth through the corona, where $\tau = 0$ corresponds to the z -axis, and $\tau = 1$ the disk surface. Between $\tau = 1$ and $\tau = 0.1$, where most of the scattering events occur, and roughly 50% of the coronal cooling takes place, the temperature is always between 3 and 20 keV. This is a relatively low temperature for a disk corona, resembling more a warm atmosphere than a hot corona. Only in the upper corona, where the density and optical depth are least, does the temperature surpass 100 keV. Yet, because it accounts for the other 50% of the coronal cooling, even this small amount of hot gas is sufficient to

Fig. 9.— Electron temperature in the corona for a converged solution of the global radiation field, for the same snapshot as in Figure 2. Within the disk photosphere, all the radiation is thermalized, and we assume the temperature is uniform for constant (r, ϕ) .



contribute a hard power-law tail to the X-ray continuum.

Another interesting feature seen in Figure 11 is the turn-over in temperature for $\tau \lesssim 10^{-3}$. As can be seen from the temperature map in Figure 9, this region is very close to the funnel/jet region, where significant outflows are expected. Since *Harm3d* cools only *bound* matter, $\mathcal{L}$ is set to zero for much of this region, leading to a decreased average temperature. This is, of course, an artifact of the simulation. Because the black hole in this simulation does not rotate, the jet power is very small and this artifact should be unimportant; when the black hole rotates and the jet power is greater, the dissipation rate in the jet could be significant. At the same time, the large relativistic bulk motion of gas in this region can still lead to interesting Comptonization effects, as will be described below in Section 6.

Finally, we comment on the relationship between our calculated electron temperature

Fig. 10.— Mean coronal temperature as a function of radius, weighted by local cooling rate $\mathcal{L}$ (top) and electron density n_e (bottom), for a range of luminosities.





Fig. 11.— Mean coronal temperature as a function of optical depth at a range of radii, for $L = 0.1L_{\text{Edd}}$.



T_e and the nominal gas temperature T found in the simulation. In a fully self-consistent picture, we would expect T_e to be very close to the ion temperature T_i because the coronal densities are generally high enough to make ion-electron collisional coupling quite rapid: just above the photosphere, the electron heating rate due to Coulomb collisions with hotter ions is

$$\frac{d}{dt} \ln(3/2kT_e) \simeq 0.14 \left(\frac{H/r}{0.06} \right)^{-1} \left(\frac{r}{10M} \right)^{-1} \frac{g(T_e)}{7} (T_i/T_e - 1) M^{-1}, \quad (12)$$

where we have estimated the local electron density by $\simeq (\sigma_T H)^{-1}$. The function $g(T_e)$ is the fitting formula derived by Mahadevan (1997); it is $\simeq 7$ for $T_e \simeq 100$ keV, increases sharply for lower temperatures, and decreases more gradually for higher temperatures. Thus, the ion-electron thermal equilibration time should be short compared to the dynamical time everywhere in the corona.

Nonetheless, we find that the ratio T_e/T is in general considerably less than unity. For example, if $\dot{m} = 0.1$, it ranges from $\sim 5 \times 10^{-3}$ (very near the photosphere) to ~ 0.2 (in the hottest locations at high altitude). For other values of $\dot{m}$, this ratio will scale like T_e because T is a fixed property of the simulation. However, as we remarked before, the gas pressure in the simulation should be interpreted as a proxy for a combination of genuine

gas pressure and radiation pressure. In this same example, the ratio of radiation pressure to nominal gas pressure in the corona varies from $\simeq 2$ –5 (in much of the coronal volume, especially near the photosphere) to $\simeq 20$ –30 (at higher altitudes and near the polar axis). This ratio varies only weakly with $\dot{m}$ because the nominal gas pressure is exactly $\propto \dot{m}$, while the general scale, but not necessarily specific local values, of the radiation pressure is likewise $\propto \dot{m}$. Thus, the majority of the nominal gas pressure should be regarded as actually due to radiation pressure, consistent with our expectation that $T_i = T_e$ and the small ratio of T_e/T . Moreover, this result suggests that a calculation in which radiation forces are included would result in a somewhat more extended corona than the one found in ThinHR.

4. BROAD-BAND SPECTRA


Having converged on a self-consistent, global map of the coronal electron temperature, there is little left to do but “turn the crank” with *Pandurata*, ray-tracing as many photon packets as computationally reasonable. As described in Schnittman & Krolik (2012), the photon packets are emitted from the photosphere of the disk with a (diluted) thermal spectrum, and subsequently up-scattered via inverse Compton in the corona, eventually either getting captured by the black hole or reaching an observer at infinity. Those photons that escape are binned by their energy and observer coordinates (θ, ϕ) , making it trivial to generate simulated X-ray spectra as a function of viewing angle. Since the *Harm3d* data is fundamentally dynamic, it is also straight-forward to simulate X-ray light curves and investigate timing features such as quasi-periodic oscillations and time lags between hard and soft bands. An in-depth study of these topics will be the subject of a future paper, but in this work, we generally average the spectra over multiple snapshots.

Specifically, we use ThinHR simulation data from snapshots between $10000M$ and $15000M$, sampled every $500M$ in time. Our photon packets cover the range of energy from 10^{-3} to 10^4 keV, with logarithmic spacing and spectral resolution of $\Delta E/E = 0.016$. By using multi-energy photon packets, we are able to resolve the thermal continuum with high accuracy and efficiency. However, the Monte Carlo IC scattering kernel still introduces a significant amount of numerical noise at high energies (Schnittman & Krolik 2012). This noise can be reduced by increasing the number of rays traced, but in practice seems to converge only slowly.

Figure 12 shows the observed spectra from ThinHR for a range of $\dot{m}$, integrated over all viewing angles. The dominant features include a broad thermal peak around 1–3 keV, a

power-law tail, and a Compton reflection hump above 10 keV.¹ There is no evidence for a cutoff up to at least 1000 keV, but the Monte Carlo statistics are very poor at those high energies. Also visible is a broadened iron line feature around 5–7 keV, which will be discussed in greater detail in the following section.

Fig. 12.— Broad-band X-ray spectra for $M = 10M_{\odot}$ and a range of luminosities, integrated over all inclination angles. In each case, the spectrum includes a broad thermal peak around 1–3 keV, a power-law tail and Compton reflection hump above 10 keV, and a broad iron line at 5–7 keV.



The most important result to be seen in Figure 12 is that, *for the first time, we have been able to use the genuine physics of global MHD simulations to reproduce the X-ray spectra observed in a wide variety of black hole binary states. Moreover, we are able to do so even while retaining an optically thick thermal disk extending to small radii.* In Table 2 we give a summary of the spectra plotted in Figure 12, using the classifications defined by Remillard & McClintock (2006). We estimate the disk fraction f^b between 2–20 keV by fitting the

¹Because we do not yet include photoionization losses other than Fe K-shell ionization, reflection at energies below 7 keV is not suppressed. At the temperatures characteristic of the inner regions of accretion disks around stellar-mass black holes, this should be a reasonable approximation, although it fails for AGN.

thermal peak with a diluted black-body spectrum from a standard N-T disk. Note that f^b is (particularly for small $\dot{m}$) smaller than $1 - L_{\text{cor}}/L_{\text{tot}}$ shown in Table 1 because it is the fraction only within the 2–20 keV band, and much of the disk power is radiated at lower energies. The power-law index Γ (where the number flux of photons per unit energy E is $N(E) \propto E^{-\Gamma}$) is measured between 10 and 100 keV. While we do not claim to completely fit the spectra with only a thermal peak and a single power-law tail, f^b and Γ are still valuable parameters for spectral classification.

Table 2: Broad-band spectral properties for a range of mass accretion rates $\dot{m}$. The disk has a peak temperature T_{disk} and contributes a fraction f^b to the total flux in the 2–20 keV band. The power-law index Γ is measured between 10 and 100 keV, and the state corresponds to the classification of Remillard & McClintock (2006).

$\dot{m}$	kT_{disk} (keV)	f^b	Γ	state
0.01	0.22	0.15	1.6	hard
0.03	0.29	0.40	2.0	hard/SPL
0.1	0.39	0.70	2.8	SPL
0.3	0.51	0.80	3.4	SPL/thermal
1.0	0.66	0.90	4.5	thermal

When scaling the simulations to $\dot{m} = 0.01$, we reproduce the low-hard state described in Remillard & McClintock (2006), with $\Gamma < 2.1$ and the 2–20 keV flux dominated by the corona: $f^b < 0.2$. At $\dot{m} = 0.1$ and above, the spectra closely resemble observations of the steep power-law (SPL) state, with $\Gamma > 2.4$ and a disk contribution of $0.2 < f^b < 0.8$. At the highest luminosities, we are safely in the thermal state, defined by $f^b > 0.75$ and little variability (see below, Sec. 7).

Despite the remarkable success of reproducing such a wide range of spectral behavior with a single simulation, we should note that these spectra represent just a one-dimensional slice through the hardness-luminosity plane that is populated by stellar-mass black holes with a wide diversity of behaviors. For example, LMC X-3 alone has been observed with $\dot{m}$ anywhere from < 0.03 up to > 0.5 *in the thermal state alone* (Steiner et al. 2010). Furthermore, with our current techniques, we are not able to reproduce the really pure thermal spectra used for inferring spin with the continuum fitting technique (McClintock et al. 2006). A sampling of the complete range of X-ray states will likely require more simulations covering different disk thicknesses, magnetic field topologies, and black hole spins.


We also note that the thermal disk component in the $\dot{m} = 0.03$ case is somewhat larger than that traditionally inferred in the low-hard state. However, more recent analyses have shown clear evidence for an optically thick thermal disk in the hard state for many galactic black holes (Miller et al. 2006a,b; Reis et al. 2010; Hiemstra et al. 2011). In those observations, the disk extends in to the ISCO with a relatively cool temperature of ~ 0.2 – 0.35 keV, consistent with what we find for $\dot{m} \leq 0.03$. Clearly the canonical disk+corona geometry for black hole accretion can extend even to the low-hard state, where a substantial fraction of thermal disk photons can get efficiently up-scattered to create a hard spectrum.

The key element in our model that leads to the shift in spectral shape with accretion rate is our assumption that the boundary of the corona is defined by the disk’s scattering photosphere. That assumption is physically reasonable: as shown in our calculation, the coronal temperature declines close to the disk surface; in addition, the quickly rising density there leads to a much greater importance for cooling processes like bremsstrahlung. Its consequence, for fixed H , is that as $\dot{m}$ increases, a larger fraction of the dissipation takes place inside the thermal disk and a smaller fraction in the corona. The disk is masked at low accretion rates because the corona, although only marginally thick by assumption, can give so much energy to the average photon created thermally in the disk. In Section 8, we will discuss how a more realistic disk picture, in which H changes with $\dot{m}$, may affect our spectral predictions.

We have also investigated the effect of observer orientation on the shape of the spectrum, but in most cases find only weak dependence on inclination. After accounting for projection effects, we do see that high-inclination (i.e., edge-on) systems have a smaller thermal peak, and somewhat harder spectrum between 1 and 10 keV, consistent with the results found in Noble et al. (2011). Above ~ 30 keV, the spectra are virtually identical. This is quite reasonable considering the results of the previous section, where we showed the coronal temperature increasing significantly with distance above the disk. The highest-energy photons are mostly generated in a large, diffuse volume in the upper corona, which subtends roughly the same solid angle independent of viewer inclination. The only part of the spectrum that seems to be strongly sensitive to the inclination is the broad iron line, which will be described in the following section.

In Figure 13 we plot the broad-band spectra for $\dot{m} = 0.01$, 0.1, and 1.0, showing the relative contributions from different regions of the accretion disk. The spectra are sorted by the emission radius of the seed photon. We see a few clear trends across all accretion rates, none of which is very surprising: the spectra grow systematically softer with increasing radius, the thermal emission is dominated by flux originating from $r/M > 15$, the coronal emission is dominated by the region $6 < r/M < 15$, and the plunging region contributes

Fig. 13.— Broad-band spectra decomposed into relative contributions from different radii in the disk. From top to bottom, $L/L_{\text{Edd}} = 0.01, 0.1, 1.0$.




very little to either the thermal or power law parts of the spectrum.

The fact that the plunging region contributes so little to the spectrum does not mean that the classical N-T disk is an adequate model for accretion dynamics. As shown in Noble et al. (2011) and Kulkarni et al. (2011), the radial emissivity profile from MHD simulations leads to thermal spectra that are systematically harder than Novikov-Thorne would predict for the same spin parameter. In part, this comes from the small amount of dissipation from inside the ISCO, but an even more important cause is the emission profile immediately outside the ISCO, which peaks at a smaller radius than predicted by Novikov-Thorne, and thus the MHD thermal spectra look like they come from BHs with somewhat higher spins (Noble et al. 2011; Kulkarni et al. 2011).

One of the major underlying assumptions of this paper is that the disk photosphere is placed at the $\tau_{\text{phot}} = 1$ surface. We believe this is an eminently reasonable and physically-motivated assumption, but it is worth investigating how much our central results are sensitive to it. To this end, we have repeated the entire radiation-temperature iterative solution for $\dot{m} = 0.1$, setting $\tau_{\text{phot}} = 0.5$ and again with $\tau_{\text{phot}} = 2.0$. The results are shown in Figure 14.

Fig. 14.— Broad-band X-ray spectra for $\dot{m} = 0.1$, when varying the optical depth of the photosphere. While the total fraction of hard X-ray flux is directly proportional to the total fraction of dissipation in the corona, the shape of the spectrum appears to be largely independent of this parameter.



While the total luminosity is of course unchanged, the relative flux in the hard X-ray tail necessarily increases with τ_{phot} . However, while the normalization of this hard tail changes, the *shape* appears to be invariant, indicative of an identical temperature profile in the upper corona, regardless of the exact location of the photosphere. This follows from the basic nature of inverse Compton radiation: equation (9) shows that the electron temperature is set by the total radiation density, not the spectrum. So the energy balance in the upper corona is completely insensitive to the detailed radiative processes taking place in the disk and boundary layer.

5. IRON EMISSION LINES

Relativistically broadened Fe $K\alpha$ lines have been detected in numerous AGN (Tanaka et al. 1995; Nandra 2007; Brenneman & Reynolds 2009), galactic black holes (Miller et

al. 2004; Reis et al. 2008, 2009), and galactic neutron stars (Cackett et al. 2010). The underlying emissivity profile is nearly always inferred (e.g., Reynolds & Nowak (2003)) by fitting the observed line profile to a phenomenological model in which the emissivity is zero inside the ISCO, rises abruptly to a maximum at the ISCO, and then declines as a power-law (sometimes a broken power-law) toward larger radii. The energy with which fluorescence photons arrive at a distant observer depends on the radius from which they are emitted and the direction in which they are sent, as well as the character of the spacetime in which they travel. It would, of course, be highly desirable both to find functional forms for the $K\alpha$ emissivity that are more closely connected to physical considerations and to be able to use observational data to constrain the disk dynamics responsible for generating these lines.

To do so requires solving problems both of physics and of procedure. Fluorescence line production begins with the illumination of gas by X-rays of energy greater than the threshold for K-shell ionization; we must determine its intensity as a function of radius. The fraction of those photons absorbed by such ionization events depends on the total optical depth of the gas and the ratio between the absorption opacity and other opacities (predominantly Compton scattering). The total optical depth depends on the specifics of angular momentum transport within the disk. The absorption opacity (as well as the fluorescence yield and the line energy) depend on the ionization state of the Fe atoms, and that in turn depends on both the temperature in the absorbing layer and the ratio between the ionization rate and the recombination rate. Although relativistic ray-tracing in vacuum has long been a solved problem (Carter 1968; Bardeen et al. 1972), a significant fraction of $K\alpha$ photons traversing a marginally optically thick corona may also gain or lose energy by Compton scattering. A significant procedural problem is posed by the question of how to separate line photons from the continuum (Miller 2007). This is particularly problematic in the case of stellar-mass black holes, where the thermal peak and the power-law tail intersect right around the iron line, making it challenging to determine the precise form of the underlying continuum spectrum.

Our new ray-tracing analysis of the *Harm3d* simulations directly solves many of these problems. The radial profile of hard X-ray illumination is a direct product of our coronal solution. The total optical depth of the disk is automatically computed by the underlying general relativistic MHD simulation, subject only to scaling with our choice of $\dot{m}$. Compton scattering *en route* also follows naturally from our Monte Carlo transfer solution. Even the continuum contribution is also an automatic byproduct, greatly improving our ability to uniquely fit the shape of the iron line.

The principal remaining uncertainty is calculation of the ionization state. In this paper we assume a fixed ionization state, but the data required for a genuine calculation of the

ionization state as a function of position are also supplied by the other components of our method, so even this last problem can be solved within our framework, although it may involve a certain amount of additional labor.

As the disk seed photons are scattered through the corona, many eventually return, with higher energy, to the disk photosphere. For stellar-mass black holes with disk temperatures of ~ 1 keV, essentially all Fe atoms will be ionized to only a few remaining electrons. However, the ability to produce a $K\alpha$ photon as a result of K-shell photoionization disappears only when the Fe is completely stripped. From Saha equilibrium, we find that most of the photosphere is dominated by He-like Fe XXV for $\dot{m} \leq 0.1$, but a mix of He-like, H-like, and fully stripped Fe exists for $\dot{m} \geq 0.3$.


The K-edge threshold varies slowly with ionization state, but for most stages is slightly above 7 keV (Kallman et al. 2004). At photon energies much above this threshold, the photoionization cross section decreases sharply with energy. In Figure 15 we show the radial profile of the absorbed K-edge photon flux, defined as the incident photon number flux in the 7–30 keV band times the fraction of the disk that is optically thick at that radius². For comparison, for $\dot{m} = 0.1$, we also show (dashed line) the number flux of seed photons emitted from the disk F_{em} with energy greater than ~ 3 keV, i.e., those most likely to get up-scattered to > 7 keV. Note that this plot of the outgoing flux is normalized for comparison purposes.

In the outer regions of the disk, the absorbed K-edge flux profile is similar in shape to the emitted flux of photons above 3 keV. However, in the plunging region the disk becomes optically thin when $\dot{m} < 0.1$ (see Table 1), suppressing the absorption there. For $\dot{m} \gtrsim 0.1$, the optically thick disk extends all the way to the horizon, and the shape of the absorbed flux profile is essentially independent of accretion rate. Even as the disk emission falls off with smaller radius, the illumination profile continues to rise as $F_{\text{Fe}} \sim r^{-\alpha}$, with $\alpha \approx 3/2$, indicative of the increasing importance of coronal flux in the inner disk.

We see no evidence for a *steepening* of the radial illumination profile with decreasing radius, as suggested by some AGN observations (Fabian et al. 1989) and predicted by “lamp post” models (Dabrowski & Lasenby 2001; Wilkins & Fabian 2011), in which the hard flux comes from a concentrated region along the black hole axis. This is not very surprising, considering the density and luminosity maps in Figures 2 and 4, which show an evacuated funnel around $\theta = 0$, essentially the opposite of the lamp post geometry. Indeed, a stationary point source on the rotation axis seems rather unlikely dynamically: a centrifugal barrier prevents much matter from approaching close to the axis, and any matter with little enough

²The reflection edge of the disk is not a sharp boundary; at a given radius, the optical depth as a function of azimuth and time can change by a factor of a few.

Fig. 15.— Absorbed iron K-edge photon flux in the local fluid frame, assuming a uniform ionization state. For lower luminosities, the disk becomes optically thin outside the horizon, leading to a clear turnover and cutoff of K-shell excitation in the plunging region. Also shown (dashed line) is the outgoing seed photon flux $F_{\text{em}}(E > 3 \text{ keV})$ for $\dot{m} = 0.1$, normalized to appear on the same axes.



angular momentum to enter that region must either fall rapidly into the black hole or be ejected; in both cases, there would be strong beaming of any photons emitted in the direction of travel.

As mentioned above, the line profile is also sensitive to the iron ionization state as a function of radius. Even if the surface density of the disk remains large inside of the ISCO, a line can be produced only if the iron is not fully ionized (Reynolds & Begelman 1997). Since we know the vertical density profile as well as the incident spectrum at each point in the disk, it should be possible to completely solve the ionization balance equations as in Garcia & Kallman (2011) and Garcia et al. (2011). Such a detailed treatment is beyond the scope of the present work, but we can substitute reasonable approximations to obtain useful first-order results.

When a photon packet hits the disk photosphere, some part is absorbed by the iron

atoms, while the remainder is reflected by electron scattering (other processes, such as free-free absorption in the disk, are insignificant). For a single photon incident on the disk, the probability of absorbing the photon in Fe K-shell photoionization is

$$P(E) = \frac{N_{\text{scat}} \kappa_{K\alpha}(E)}{N_{\text{scat}} \kappa_{K\alpha}(E) + \kappa_T}, \quad (13)$$

where $\kappa_{K\alpha}$ and κ_T are the Fe $K\alpha$ and Thomson scattering opacities. N_{scat} is the median number of scattering events a photon experiences before emerging from the atmosphere. Thus, the typical photon gets N_{scat} chances to excite a $K\alpha$ transition before exiting the disk, thereby enhancing the yield on the line production (Kallman et al. 2004). For accretion disks with roughly solar abundances, we take

$$\begin{aligned} \kappa_{K\alpha}(E < 7 \text{ keV}) &= 0 \\ \kappa_{K\alpha}(E > 7 \text{ keV}) &= \kappa_T \left(\frac{E}{7 \text{ keV}} \right)^{-3} \end{aligned} \quad (14)$$

and from Monte Carlo scattering experiments, we find an angle-averaged value of $N_{\text{scat}} = 3$. This crude approximation to the K-shell opacity is appropriate provided most Fe atoms retain at least one electron; when most Fe atoms are stripped, $\kappa_{K\alpha}$ is smaller than our estimate by the ratio of FeXXVI ions to the total. Of all the photons absorbed by iron in the disk, only a fraction $f_{K\alpha}$ produce a fluorescent line, while the excitation energy deposited by the rest is lost to Auger transitions, or, in the case of H-like and He-like Fe, more energetic K series recombination lines (Kallman et al. 2004). In Pandurata’s current form, the energy absorbed by K-edge opacity is simply removed from the spectrum during the Monte Carlo solution, while the energy in $K\alpha$ emission is added back later. The fluorescence yield $f_{K\alpha}$ depends on ionization state (Krolik & Kallman 1987), growing slowly from $\simeq 0.34$ to $\simeq 0.5$ from FeI to FeXXII. At higher ionization stages, it can be as little as 0.11 (FeXXIII), but is generally larger (0.5–0.75). For all the results presented below, we take $f_{K\alpha} = 0.5$, corresponding to a highly ionized state.

For a photon packet incident on the disk with initial spectral intensity $I_{\nu,0}$ [units of erg/s/Hz], the number of Fe $K\alpha$ photons produced per second will be

$$N_{K\alpha} = f_{K\alpha} \int d\nu P(h\nu) \frac{I_{\nu,0}}{h\nu}. \quad (15)$$

In our simplified model, all of these photons are added back to the photon packet as a delta function in energy at $E = 6.4 \text{ keV}$ ³. Including absorption, the outgoing spectrum can be

³The threshold and emission energies are taken from neutral iron, which is technically more appropriate for AGN applications. For the highly ionized Fe expected in galactic black hole sources, the lines will be emitted at slightly higher energies, up to 6.9 keV for FeXXVI.


written

$$I_{\nu,\text{out}} = I_{\nu,0}[1 - P(h\nu)] + N_{K\alpha}\delta(h\nu - 6.4 \text{ keV}) \cdot (6.4 \text{ keV}), \quad (16)$$

where both $I_{\nu,\text{out}}$ and $I_{\nu,0}$ are measured in the local frame of the disk.

This outgoing photon packet then propagates towards the observer, getting up-scattered by coronal electrons, possibly returning to the disk, or captured by the black hole. The Fe $K\alpha$ emission line and absorption edge are both broadened by relativistic effects as well as the IC scattering. To show the magnitude and shape of these spectral features, in Figure 16 we plot (black curve) the ratio of the total observed spectrum to what would be observed if no absorption or emission were included in the calculation. We also show the absorption and emission contributions separately with the red and blue curves, respectively.

Fig. 16.— Iron line absorption and emission features as measured by an observer at infinity, including all relativistic effects, for $\dot{m} = 0.1$ and $i = 45^\circ$. The curves show the ratio of the observed flux to that which would be observed without line physics included. The red curve shows only the absorption edge, the blue curve shows only the emission line, and the black curve shows the combination of absorption and emission. Above 10 keV, the sharp spectral features are due to Monte Carlo noise.



When plotted as a ratio to the hypothetical no-line spectrum, the absorption appears

to dominate over the emission. Two facts account for this effect. First, $f_{K\alpha} = 0.5$ means that twice as many photons get absorbed as emitted. Second, although in the rest-frame only photons with energy > 7 keV can be absorbed and all line photons have energy exactly 6.4 keV, relativistic broadening can shift part of the absorption feature downward in energy and part of the emission line upward. Where they overlap, there is substantial cancellation. Bluewards of the point where they exactly cancel, the spectrum shows a sharp absorption feature that looks like $1 - P(E)$. Of course, the line ratios in Figure 16 could never be directly observed, since they require knowledge of some hypothetical spectrum that conveniently ignores fluorescent line physics, but they do provide valuable insight into the physical processes at work here.


In practice, we observe spectra like those shown in Figure 12 and then attempt to infer the shape of the emission line by fitting the continuum with phenomenological models, an approach that can introduce serious systematic errors (e.g., Miller (2007) and references therein). The great advantage of this global radiation transport calculation is that *we can simultaneously fit the entire spectrum with a single model based on physical parameters*—black hole mass, spin, and accretion rate, Fe abundance, and observer inclination angle, obviating the historical reliance on more phenomenological models. Ultimately, we hope to apply our global radiation transport techniques to a large body of MHD simulations, resulting in a comprehensive suite of tabulated, self-consistent spectra that can be incorporated into a standard X-ray spectra analysis package like XSPEC (Arnaud 1996). As mentioned above, more detailed ionization physics will be required before *Pandurata* can be used to fit real iron line data with high precision. Nonetheless, in the meantime we can still use the simulated spectra to gain important insights into the behavior of the inner accretion flow.

In Figure 17 we plot the shape of the iron line (ratio of “emission only” to “no line physics”) for a range of observer inclination angles for $\dot{m} = 0.01$ (the thinner disk highlights the relativistic effects). For these lines, we find equivalent widths of ≈ 200 –350 eV, consistent with observations of stellar-mass black holes as well as AGN (Miller et al. 2004, 2006a; Walton et al. 2012). The basic shapes of the lines also agree closely with the thin-disk calculations well-known in the literature for over two decades (e.g., Laor (1991)).

One important difference that is not often discussed is the high-energy tail clearly seen at all inclinations (Petrucci et al. 2001). This is due to IC scattering in the corona, the very process that generates the ionizing flux in the first place. Since the coronal scattering is nearly isotropic, the amplitude of the emission line above ~ 8 keV is independent of viewer inclination. At the same time, the optical depth from the disk directly to a distant observer increases with inclination angle, thereby reducing the total number of line photons that *don’t* get up-scattered, as can be seen by comparing the integrated line flux below 8 keV.

The diminished contrast between line and continuum associated with high inclination angles may make it systematically more difficult to detect $K\alpha$ lines from those directions.


Fig. 17.— Iron line profile as a function of observer inclination, for $\dot{m} = 0.01$. Only the emission contribution is shown. The extended blue tail above 8 keV is due to inverse-Compton scattering of the line photons in the corona.



One of the most important potential applications of the iron line is to use it as a probe of where (or if) the disk truncates. If there is a truncation radius, and this radius can be quantitatively related to the ISCO, measurement of the truncation radius could lead to a measurement of the black hole spin. Indeed, many have attempted to measure spin *assuming* that such a sharp truncation takes place exactly at the ISCO (a sampling of these efforts may be found in Martocchia et al. (2002); Miller et al. (2002); Duro et al. (2011); Reis et al. (2011, 2012); Fabian et al. (2012)). Even for the single spin value ($a/M = 0$) simulated in ThinHR, we have shown in Table 1 and Figure 15 that the reflection edge of the disk can be adjusted by modifying $\dot{m}$. Therefore we might reasonably expect very different line profiles for $\dot{m} = 0.01$, 0.03, and 0.1, corresponding to average reflection edge radii $R_{\text{refl}}/M = 6.1$, 4.4, and 2.1, respectively. Gravitational redshift is especially strong in the plunging region, so sizable contrasts in the red portion of the profile might be expected. A sample of line profiles is shown in Figure 18 for a range of $\dot{m}$, holding the observer inclination constant at

$i = 30^\circ$. Remarkably, these lines show essentially no variation with R_{refl} , especially in the red wing below 6 keV.

Fig. 18.— Iron line profile as a function of luminosity, for observer inclination $i = 30^\circ$. The red wings of the lines are remarkably similar, despite differences in R_{refl} and illumination profiles.



The only differences evident are the height of the peak and the amplitude of the IC tail. Both of these features may be attributed to the different effective geometries, as the higher luminosity cases have a greater scale height for the disk and corona, somewhat increasing the optical depth through which a line photon must pass before reaching the observer. In the limit where the observer inclination angle is the same as the disk opening angle, the optical depth through a sandwich-type corona becomes infinite, and it becomes harder for line photons to escape directly to the observer without scattering. This finite-opening-angle effect also explains the different shape of the line for $\dot{m} = 1.0$. With a photospheric aspect ratio $H_{\text{phot}}/r \approx 0.5$, its configuration can hardly be called a thin disk. In that case, an observer at 30° would actually be sampling a range of effective inclinations between 0° and 60° , and thus the total integrated line looks like a combination of these different viewing angles.

This lack of sensitivity to the inner disk location is due to the fact that gas inside of the ISCO is already plunging rapidly towards the horizon. Most of the line photons produced in the plunging region get beamed into the black hole, never reaching a distant observer. Similar results were seen for thermal emission from the plunging region in Zhu et al. (2012). As a test of this effect, we compared the total flux that eventually reached infinity with that which was captured by the horizon as a function of the radius of the initial seed photons. For $\dot{m} = 0.1$, 40% of the seed photons emitted from $r = 5M$ got captured by the BH, a fraction that climbs to 95% at $r = 3M$. As can be seen in Figure 13, $\lesssim 1\%$ of the total flux around 6 keV comes from inside of $6M$, regardless of where the disk reflection edge is. In future work making use of `Harm3d` simulations for a range of spin parameters, we will explore whether the emergent profiles have sufficiently strong dependence on spin that this diagnostic can be successfully used.


6. BULK COMPTONIZATION

In addition to the thermal IC processes described above, the corona can also transfer energy to the seed photons through “bulk Comptonization” when the fluid velocity of the corona is large relative to the disk. Some authors have used this process to explain the hard tail seen in some thermal state observations (Zhu et al. 2012), or the steep power-law state when the bulk flow is convergent (Titarchuk & Shrader 2002; Turolla et al. 2002) or turbulent (Socrates et al. 2004; Socrates 2010). To quantify this effect in the `Harm3d` simulations, we simply set the electron temperature everywhere in the corona to zero while maintaining the turbulent motion above the disk and the convergent flow in the plunging region.

As before, we calculate the total Compton power in each fluid element by subtracting the energy in the incoming photon packet (as measured at infinity) from that of the outgoing ray. For an electron at rest, the photon will always transfer energy to the electron, giving negative IC power (hence the distinction between Compton scattering and *inverse* Compton scattering). In Figure 19 we plot the bulk coronal power in terms of $dL/d(\log r)$, normalized by the total disk luminosity, for a range of accretion states. The y -axis is logarithmic and signed, so we set $10^{-4} = 0 = -10^{-4}$ for improved visualization. Where $dL/d(\log r) > 0$, the bulk velocity of the gas transfers energy to the radiation field. Where $dL/d(\log r) < 0$, the typical energy of a seed photon is greater than the bulk kinetic energy of the coronal electrons, and the radiation field loses energy to the corona.

Three conclusions may be drawn from Figure 19: (1) bulk Comptonization plays a very minor role in the overall energetics of thin accretion disks; (2) bulk Comptonization is most significant for high accretion rates; and (3) for all values of $\dot{m}$, $dL/dr > 0$ in the inner regions

Fig. 19.— Net Compton power $dL/d(\log r)$ in the corona when $T_e = 0$, normalized to the total seed luminosity from the thermal disk. Note the unusual labeling of the y -axis, with logarithmic scaling above and below $0 = \pm 10^{-4}$.




and $dL/dr < 0$ in the outer disk. The explanations of these effects are straightforward: (1) turbulent velocities in the corona are simply not very large. In the plunging region, there are either not enough seeds (low $\dot{m}$) or the optical depth is high (large $\dot{m}$), so the seeds are advected into the black hole without being able to sample a wide range of converging velocities. Not surprisingly, the observed spectra for the bulk Comptonization runs are nearly indistinguishable from pure thermal disks. (2) Large $\dot{m}$ corresponds to large H_{phot}/r , and we find that the turbulent velocities generally increase with scale height above the disk, so higher luminosity systems are sampling more turbulent regions of the corona. (3) For disks with constant H_{phot}/r , turbulent velocity should scale like the orbital velocity $v_{\text{bulk}} \sim v_{\text{orb}} \sim r^{-1/2}$, so the turbulent kinetic energy scales like r^{-1} . The seed photon energy, on the other hand, scales like $r^{-3/4}$ in the outer disk, and actually begins to decrease in the inner region as the disk becomes optically thin.

In Figure 20 we show the average kinetic energy in the corona as a function of radius (solid curves), along with the seed photon energy (dashed curves). Here, the specific kinetic energy is defined as $1/2 \sigma_v^2(r)$, where $\sigma_v(r)$ is the variance of the 3-velocity u^i , sampled over

all ϕ , t , and optical depth $\tau = 0.1 - 1$ at each radius. Outside of $r \approx 10M$, the photon energy is higher, and thus transfers energy into the corona, giving $dL/dr < 0$. Note that, for $\dot{m} = 0.01$, the nearly laminar flow at the midplane is considered part of the corona, not the disk, thus explaining the turnover in turbulent kinetic energy inside of $\sim 4M$.

Fig. 20.— Seed photon energy (dashed curves) and specific turbulent kinetic energy of the fluid in the corona (solid curves) for $L/L_{\text{Edd}} = 0.01, 0.1, 1.0$.




7. X-RAY VARIABILITY

Up to this point, all the discussion in this paper has focused on steady-state behavior of the simulated spectra. As mentioned in section 2, the results are based on snapshots of the ThinHR simulation, spaced every $500M$ between $10000M$ and $15000M$, roughly the period of inflow equilibrium. With these 11 snapshots, we are also able to carry out some very coarse timing analysis. Figure 21 shows simulated light curves in four different energy bands, for $\dot{m} = 0.1$ and observer inclination $i = 60^\circ$. Over the period shown, the bolometric luminosity of the simulation changes by about 20%, quite typical of global MHD simulations. To focus on the intrinsic variability, we have normalized all light curves by a single linear trend over this period. In Figure 21, each individual light curve has also been normalized

by its mean value, to show the relative amplitude of fluctuations.

Fig. 21.— X-ray light curves for $L/L_{\text{Edd}} = 0.1$ and viewer inclination of $i = 60^\circ$, taken from simulation snapshots sampled every $500M$ in time during the period of inflow equilibrium. In each energy band, the light curve is normalized to the mean flux in that band, and divided by the linear trend of the bolometric flux.



In Table 3 we list the RMS variation in the flux in different energy bands for different values of $\dot{m}$, again normalized by the linear trend in the bolometric luminosity. At low energies, corresponding to the thermal peak, we see a clear anti-correlation between accretion rate and variability: due to the fact that at low $\dot{m}$, the inner disk is moving in and out, changing the thermal seed flux. The low variability of the case with $\dot{m} = 1.0$ further strengthens its classification in the thermal state (Remillard & McClintock 2006).

Our findings show an opposite trend in variability with $\dot{m}$ than found in Noble & Krolik (2009). They measured the bolometric flux from dissipation in the corone only. Since they neglected the seed flux from the thick disk and ignored all IC scattering/heating physics in the corona, their primary source of variability was entirely from the intrinsic dynamic variability of coronal turbulence and dissipation. Thus their trend with $\dot{m}$ suggests more inherent fluctuations higher above the disk, consistent with increased turbulent velocities,

as shown above in Figure 20.

At all $\dot{m}$, there is a clear increase in variability with photon energy, as seen in observations of the steep power-law state of some black holes (Cui et al. 1999). This could result if the high-energy flux comes from a power law with variable index, where the power-law tail “pivots” around the thermal peak like a see-saw. However, it is equally clear that the light curves in the different bands are not tightly correlated, as would be expected if the variability were strictly due to global coronal properties like the Compton y -parameter. That lack of correlation indicates that the dependence of variability on photon energy is due to fluctuations that are independent in regions of different temperature, as well as stronger in regions of higher temperature. Such a situation is a natural result of variability in local heating.

Table 3: RMS variability in different energy bands as a function of luminosity, for an observer at $i = 60^\circ$. To remove the secular trend, the light curves in each energy band have been normalized by dividing out a linear fit to the bolometric luminosity.

L/L_{Edd}	0.1-3 keV (%)	3-10 keV (%)	10-30 keV (%)	> 30 keV (%)
0.01	15	24	35	54
0.03	14	26	40	56
0.1	12	24	32	58
0.3	8.4	21	26	54
1.0	6.8	19	28	51

To get a better estimate of the local coronal effects on the light curves, we have also calculated phase-dependent light curves by calculating the flux seen by observers at different azimuths. The variability as a function of observer ϕ is a proxy for continuum fluctuations at high frequencies, comparable to the orbital frequency, where variability is often quite strong in the hard and steep power-law states (Remillard & McClintock 2006). To estimate the amplitude of these modulations, we construct many short light curves, one for each snapshot and inclination angle, and calculate the fractional RMS amplitude for each light curve. We then average over all these snapshots, plotting the mean RMS (along with 1σ error bars) in Figure 22 as a function of observer inclination for a range of energy bands. This procedure is roughly equivalent to measuring the variance in the observed light curves over a narrow frequency band corresponding to the orbital frequencies of the parts of the disk that contribute the most power. In order to resolve the high-energy fluctuations, we use a particularly large number of Monte Carlo photon packets, roughly 10^9 rays per snapshot.

8. COMPARISON WITH CLASSICAL DISK THEORY

The spectral states of galactic black hole binaries are roughly correlated with their bolometric luminosities in the sense that low luminosity states generally have hard spectra, while higher luminosities permit a broader range of spectral states but exhibit a preference for softer states (Fender et al. 2004; Remillard & McClintock 2006). Our model is able to reproduce this observed correlation, yet does so in a fashion that differs in several respects from classical disk theory. In our model, increased accretion rate leads to a proportionately larger surface density, but leaves the scale height unchanged. The systematic shift in spectral shape with accretion rate is due to a change in how the corona and the thermal disk share the dissipation: as the accretion rate increases, the photosphere moves to larger multiples of H , so that more of the dissipation occurs within the thermal disk. By contrast, in classical disk theory, both Σ and H are functions of accretion rate (Shakura & Sunyaev 1973). As we will show in a moment, most of our parameter space lies in the radiation-dominated regime, in which $\Sigma \propto \dot{m}^{-1}$ and $H \propto \dot{m}$. Moreover, for the accretion rates we consider, classical disk theory assumes that *all* the dissipation takes place inside the disk, leaving no room for a corona at all, and identifies the photosphere precisely with H . These contrasts raise two questions: Is there any $\dot{m}$ at which our model and classical theory overlap? And should the trends with $\dot{m}$ predicted by the single-simulation model presented here be expected to carry over to global MHD simulations with different values of H , corresponding to different values of $\dot{m}$?

As shown in Figure 13, most of the light is produced in the range $r \simeq 6\text{--}30M$ for all the accretion rates we studied. Our account of the predictions of conventional disk theory therefore centers on that range. According to this theory, radiation pressure exceeds gas pressure inside disks when the accretion rate is greater than

$$\dot{m}_{rg} \simeq 0.02 \alpha_{SS}^{-1/8} (M/M_\odot)^{-1/8} (r/10M)^{21/16} (R_R/0.2)^{-9/8} R_T^{1/8} R_z^{5/8}, \quad (17)$$

where α_{SS} is the usual ratio of vertically-integrated fluid-frame stress to vertically-integrated pressure, $R_R \leq 1$ is a function of radius that adjusts the vertically-integrated fluid-frame dissipation rate for both the net angular momentum flux through the disk and relativistic corrections, $R_T \simeq R_R$ is a similar correction factor applied to the vertically-integrated fluid-frame stress, and R_z (usually slightly greater than unity) introduces relativistic corrections into the vertical component of gravity (notation as in Krolik 1999). Because R_R increases outward in this range of radii, $\dot{m}_{rg}$ rises only gradually with radius. Thus, almost all the span of accretion rates we consider falls into the radiation-dominated regime.

Conventional analytic disk theory estimates disk thickness in the radiation-dominated limit by supposing that radiation provides all the disk’s support against the vertical compo-

ment of gravity and that all dissipated energy is conveyed outward by radiation flux (Shakura & Sunyaev 1973; Novikov & Thorne 1973). This pair of assumptions combined with the condition of hydrostatic equilibrium leads to the conclusion that all of the dissipation within the disk must be accomplished within a distance

$$H_r = (3/2)(\dot{m}/\eta) \left(\frac{GM}{c^2} \right) \frac{R_R}{R_z} \quad (18)$$

of the disk midplane. Thus, for radiation-dominated disks, scale height and accretion rate are directly proportional, and one can be used as a proxy for the other. Conventional theory also assumes that the density is constant for $|z| \leq H_r$ and zero outside H_r (Shakura & Sunyaev 1973). Because this theory has no explicit place for coronae, it is unclear what luminosity it predicts for them, but presumably any corona begins outside H_r .

Detailed stratified shearing box simulations of radiation-dominated disk segments (Hirose et al. 2009a; Blaes, Hirose & Krolik 2011) have shown that H_r does give an order of magnitude estimator of the vertical scale height of such disks, but, unsurprisingly, the density distribution is crudely exponential. Because both gas and magnetic pressure can contribute to vertical support and some vertical energy transport is by radiation advection rather than diffusive flux, these simulations also find that the dissipation profile is more extended than indicated by conventional analytic theory. In particular, only about half the dissipation takes place within a distance H_r of the midplane, and 90% of the dissipation is accomplished inside $\simeq 2H_r$. The photosphere is generally found at $3\text{--}4H_r$. These simulations were conducted with surface densities $\sim 10^4\text{--}10^5 \text{ gm cm}^{-2}$, a range relevant to the larger radius and smaller accretion rate portion of our parameter space, so the specific numbers quoted might require adjustment for smaller radii and larger accretion rates.

The parallel set of facts about our simulated disk is that its density scale height is always $H \simeq 0.06r$, but the height of its photosphere is a multiple of H that increases with $\dot{m}$, ranging from $\simeq 2$ to $\simeq 9$ as $\dot{m}$ increases from 0.01 to 1.0 (see Table 1). This is, of course, why the corona’s share of the luminosity decreases with increasing $\dot{m}$ in our model.

Having placed the two pictures side-by-side, we can now locate the parameters for which they resemble each other. Because $H_r \propto R_R$ and R_R scales with r somewhat more slowly than linearly for Schwarzschild spacetimes between $r = 10M$ and $r = 30M$, while the simulation has $H \propto r$, it is possible to match the simulation photosphere to the photosphere predicted by classical disk theory across this range of radii. We find approximate agreement when $\dot{m} \simeq 0.2$.

We now turn to the question of what to expect from simulations with different $H(r)$, representing other values of $\dot{m}$. Because the central issue is the relative share of dissipation

taking place inside and outside the photosphere, this question can be posed in terms of how the two ratios H_{phot}/H and H_{mag}/H (here H_{mag} is the *magnetic* scale height) depend on accretion rate. If, as in our single-simulation model, the former increases more rapidly with $\dot{m}$ than the latter, our predictions will be (at least qualitatively) vindicated; if not, they will need revision.

9. DISCUSSION

Using the new radiation transport code *Pandurata*, we have analyzed data from a high-resolution, 3D MHD simulation of accretion onto a Schwarzschild black hole. Because the MHD code *Harm3d* is energy conserving, we are able to employ a cooling function that tracks the local dissipation of energy throughout the simulation volume. By combining the results from *Pandurata* and *Harm3d*, we have for the first time been able to produce a global, self-consistent solution for the radiation field around an accreting black hole, and predicted—on the basis of real physics—the coronal luminosity. The major results from this work can be summarized as follows:

1. For different values of the Eddington-normalized accretion rate $\dot{m}$, the location of the photosphere changes, in turn varying the fraction of radiative power in the disk (thermal) and the corona (inverse Compton). The coronal temperature ranges from about 10 keV near the disk surface up to $\sim 100\text{--}300$ keV in the upper, low-density corona. Independent of black hole mass, the corona temperature scales inversely with accretion rate: $T_e \sim \dot{m}^{-1}$ at temperatures well below $m_e c^2/k_B$ and $T_e \sim \dot{m}^{-1/2}$ in the relativistic regime.
2. By varying $\dot{m}$ from 0.01 to 1, we naturally reproduce X-ray spectra consistent with those observed in the hard, steep power-law, and thermal states of galactic black hole binaries. The spectra are characterized by a thermal peak around 1 keV and a high-energy power-law tail extending above 100 keV. In most cases there is evidence for a Compton reflection hump between 30 keV and 100 keV.
3. The Fe $K\alpha$ illumination profile of the disk follows the classical r^{-3} scaling at large radius, then flattens to $r^{-3/2}$ in the inner disk. At lower values of $\dot{m}$, the disk begins to disappear inside the ISCO, and the line production is correspondingly reduced. The iron line profiles consist of both an absorption edge above 7 keV, and a broad emission line around 6.4 keV with a strong red-shifted tail. The shape of the line is dependent on observer inclination, and in all cases has a significant tail above 8 keV due to up-scattering of the line photons in the corona.

The observed iron line profile appears to be only weakly dependent on accretion rate or the location of the disk inner edge, as most photons generated in the plunging region never reach the observer. Thus iron lines may in fact be better at measuring the location of the ISCO than the disk’s reflection edge.

4. Bulk Comptonization plays a very minor role in the photon energetics, typically $\lesssim 1\%$ of the total seed luminosity. In the outer region of the disk ($r \gtrsim 10M$), the thermal seed photons carry more energy than the bulk kinetic energy in the coronal electrons.
5. We have carried out some initial timing analysis of the simulated X-ray spectra, and find a number of trends that are consistent with observations: the fractional RMS amplitude increases with decreasing luminosity, and for all accretion rates, the RMS amplitude increases with photon energy. On short time scales, the variability increases with observer inclination and photon energy, as expected for a coronal hot-spot model of X-ray variability.

Although the progress made to date has been significant, this work is just the tip of the iceberg. We are currently in the process of analyzing new *Harm3d* simulations, carried out with resolution comparable to that of *ThinHR*, for a wide range of black hole spin parameters. This will allow us to explore both potential dependence of the disk/corona continuum spectrum on spin and greatly improve our understanding of the iron line as a probe of black hole spin, disk dynamics near the ISCO, and the nature of the plunging region. New simulations with high spin will also allow us to probe the properties of the relativistic jet, frequently seen in observations of the hard state, and most clearly present in simulations of spinning black holes (McKinney & Gammie 2004; Hawley & Krolik 2006).

In other future work, simulations of disks with different $H(r)$ profiles will expand the applicability of our models to a wider range of X-ray states and accretion rates. We will explore what parameters other than $\dot{m}$ and $H(r)$ determine the state of a given black hole. Improving the physics of the fluorescent line, including ionization balance and more detailed excitation cross sections (Garcia & Kallman 2011; Garcia et al. 2011) will make predictions of its strength and profile more reliable. Including the energy lost to photoionization and Compton recoil in the disk surface will also permit a better treatment of hard X-ray reprocessing and that energy’s reemergence in disk continuum, processes relevant to AGN.

We also plan on extending our preliminary variability analysis to the entire set of simulation data within *ThinHR*’s statistically steady epoch (over 5000 snapshots in time), allowing for a more detailed study of high-frequency fluctuations and the possible identification of quasi-periodic oscillations (QPOs), sometimes seen in galactic black holes in the steep power-law state (Remillard & McClintock 2006). In addition to QPOs, we should be able to

characterize time lags between different energy bands as a function of frequency, and compare with a large body of observational results (e.g., Nowak et al. (1999)). These lags appear to scale like the light-crossing time for fluctuations in the thermal seed flux to propagate through the corona, and thus could be a powerful probe of the coronal geometry (Uttley et al. 2011). We can also investigate the effects of finite light-travel time through the simulation volume, which was found to suppress variability in Noble & Krolik (2009).

Pandurata was originally developed to study X-ray polarization (Schnittman & Krolik 2009, 2010), so that information comes along for free with all the calculations described in this paper. As techniques for high-sensitivity X-ray polarimetry continue to improve (Black et al. 2010), polarization predictions will become observationally testable; it will be interesting to compare predictions made from MHD simulations with the toy coronal models presented in Schnittman & Krolik (2010). Because our analysis includes broad-band spectra, line profiles, timing, and polarization information in a single self-consistent calculation, it is the ideal tool for integrating these complementary techniques for measuring black hole spin and probing the physical properties of the accretion flow.

Despite the remarkable progress we have made in bridging the gap between simulation and observation, there still exist numerous challenges and caveats. The state-of-the-art MHD simulations still do not include adequate thermodynamics or internal radiation transport coupled directly with the fluid dynamics. While much progress has been made in shearing-box simulations (Hirose et al. 2006, 2009a), there remain serious conceptual and computational obstacles to incorporating these advances into global simulations. The ray-tracing tools described here are also lacking in certain regards. Because of the photon-packet methodology used in **Pandurata**, we have been forced to use energy-independent scattering cross sections, which certainly breaks down at high photon energy. Similarly, the angular distribution of Fe K α lines emerging from the disk is assumed to be identical with the angular distribution of photons in the same packet reflected by Compton scattering as given by Chandrasekhar (1960). These shortcomings can be improved with relatively little effort, but at a cost to computational efficiency.

Without a doubt, the most important next step is the direct comparison of our **Pandurata** spectra with real X-ray data. To this end, we are fortunately blessed with a mass of archival data from *RXTE*, *Chandra*, *XMM-Newton*, and *Suzaku* with which to test our spectral models and improve upon earlier phenomenological analysis methods.

We would like to thank C. Done, A. Fabian, T. Kallman, and C. Reynolds for helpful discussions. This work was partially supported by NSF grants AST-0507455 and AST-0908336 (JHK) and AST-1028087 (SCN). The ThinHR simulation was carried out on the

Teragrid Ranger system at the Texas Advance Computing Center, which is supported in part by the National Science Foundation.

REFERENCES

- Arnaud, K. A. 1996, in *Astronomical Data Analysis Software and Systems V*, ASP Conference Series, 101, 17
- Balbus, S. A., & Hawley, J. F. 1991, ApJ, 376, 214
- Bardeen, J. M., Press, W. H., & Teukolsky, S. A. 1972, ApJ 178, 347
- Beckwith, K., Hawley, J. F., & Krolik, J. H. 2008, MNRAS, 390, 21
- Elack, J.K. et al., 2010, Proc. SPIE, 7732, 77320X-1
- Blaes, O.M., Hirose, S. & Krolik, J.H. 2011, ApJ, 733, 110
- Brenneman, L. W., & Reynolds, C. S. 2009, ApJ, 702, 1367
- Cackett, E. M., et al. 2010, ApJ, 720, 205
- Carter, B. 1968, Phys. Rev., 174, 1559
- Chandrasekhar, S. 1960, *Radiative Transfer*, Dover, New York
- Cui, W., Zhang, S. N., Chen, W., & Morgan, E. H. 1999, ApJ, 512, L43
- Dabrowski, Y., & Lasenby, A. N. 2001, MNRAS, 321, 605
- Done, C. & Kubota, A. 2006, MNRAS, 371, 1216
- Duro, R. et al. 2011, A&A, 533, L3
- Elvis, M., et al. 1978, MNRAS, 183, 129
- Fabian, A. C., Rees, M. J., Stella, L., & White, N. E. 1989, MNRAS, 238, 729
- Fabian, A.C. et al. 2012, MNRAS
- Fender, R. P., Belloni, T. M., & Gallo, E. 2004, MNRAS, 355, 1105
- Garcia, J., & Kallman, T. R. 2011, ApJ, 718, 695
- Garcia, J., Kallman, T. R., & Mushotzky, R. F. 2011, ApJ, 731, 131
- Haardt, F. & Maraschi, L. 1993, ApJ, 413, 507
- Haardt, F., Maraschi, L., & Ghisellini, G. 1994, ApJ, 432, L95

- Hawley, J. F., & Krolik, J. H. 2006, *ApJ*, 641, 103
- Hawley, J. F., Guan, X., & Krolik, J. H. 2011, *ApJ* 738, 84
- Hiemstra, B., Mendez, M., Done, C., Diaz Trigo, M., Altamirano, D., & Casella, P. 2011, *MNRAS*, 411, 137
- Hirose, S., Krolik, J. H., & Stone, J. M. 2006, *ApJ*, 640, 901
- Hirose, S., Krolik, J. H., & Blaes, O. 2009a, *ApJ*, 691, 16
- Hirose, S., Blaes, O., & Krolik, J. H. 2009b, *ApJ*, 704, 781
- Kallman, T. R., Palmeri, P., Bautista, M. A., Mendoza, C., & Krolik, J. H. 2004, *ApJS*, 155, 675
- Krolik, J. H. & Hawley, J. F. 2002, *ApJ*, 573, 754
- Krolik, J. H. & Kallman, T. R. 1987, *ApJ*, 320, L5
- Kulkarni, A. K., et al. 2011, *MNRAS*, 414, 1183
- Laor, A. 1991, *ApJ*, 376, 90
- Liang, E.P.T. & Price, R.H. 1978, *ApJ*, 218, 247
- Mahadevan, R. 1997, unpublished Harvard Ph.D. thesis
- Martocchia, A., Matt, G., Karas, V., Belloni, T. & Feroci, M. 2002, *A&A*, 387, 215
- McClintock, J. E., Shafee, R., Narayan, R., Remillard, R. A., Davis, S. W., Li, L.-X. 2006, *ApJ*, 652, 518
- McKinney, J. C., & Gammie, C. F. 2004, *ApJ*, 611, 977
- Miller, J.M. et al. 2002, *ApJ Letts*, 570, L69
- Miller, J. M., et al. 2004, *ApJ*, 606, L131
- Miller, J. M., Homan, J., & Miniutti, G. 2006a, *ApJ*, 652, L113
- Miller, J. M., et al. 2006b, *ApJ*, 653, 525
- Miller, J. M. 2007, *ARA&A*, 45, 441

- Mitsuda, K., Inoue, H., Koyama, K., Makishima, K., Matsuoka, M., Ogawara, Y., Suzuki, K., Tanaka, Y., Shibasaki, N., & Hirano, T. 1984, PASJ, 36, 741
- Nandra, K., O'Neill, P.M., George, I.M. & Reeves, J.N. 2007, MNRAS 382, 194
- Noble, S. C., Krolik, J. H., & Hawley, J. F. 2009, ApJ 692, 411
- Noble, S. C., & Krolik, J. H. 2009, ApJ 703, 964
- Noble, S. C., Krolik, J. H., & Hawley, J. F. 2010, ApJ 711, 959
- Noble, S. C., Krolik, J. H., Schnittman, J. S., & Hawley, J. F. 2011, ApJ 743, 115
- Novikov, I. D., & Thorne, K. S. 1973. in *Black Holes*, ed. C. DeWitt & B. S. DeWitt (New York: Gordon and Breach)
- Nowak, M. A., Wilms, J., & Dove, J. B. 1999, ApJ, 517, 355
- Oda, M., Gorenstein, P., Gursky, H., Kellogg, E., Schreier, E., Tananbaum, H., & Giacconi, R. 1971, ApJ, 166, L1
- Petrucchi, P. O., Merloni, A., Fabian, A., Haardt, F., & Gallo, E. 2001, MNRAS, 328, 501
- Poutanen, J., Krolik, J. H., & Ryde, F. 1997, MNRAS, 292, L21
- Reis, R. C., Fabian, A. C., Ross, R. R., Miniutti, G., Miller, J. M., & Reynolds, C. S. 2008, MNRAS, 387, 1489
- Reis, R. C., Fabian, A. C., Ross, R. R., & Miller, J. M. 2009, MNRAS, 395, 1257
- Reis, R. C., Fabian, A. C., & Miller, J. M. 2010, MNRAS, 402, 836
- Reis, R.C. et al. 2011, MNRAS, 410, 2497
- Reis, R.C., Miller, J.M., Reynolds, M.T., Fabian, A.C. & Walton, D.J. 2012, ApJ, 751, 34
- Remillard, R. A., & McClintock, J. E. 2006, ARA& A, 44, 49
- Reynolds, C. S., & Begelman, M. C. 1997, ApJ, 488, 109
- Reynolds, C. S., & Nowak, M. A. 2003, Phys. Reports, 377, 389
- Rybicki, G. B., & Lightman, A. P. 2004, *Radiative Processes in Astrophysics* (Weinheim: Wiley-VCH)
- Sano, T., Inutsuka, S., Turner, N.J. & Stone, J.M. 2004, ApJ, 605, 321

- Schnittman, J. D. 2005, *ApJ*, 621, 940
- Schnittman, J. D., & Rezzolla, L. 2006, *ApJ*, 637, L113
- Schnittman, J. D., Krolik, J. H., & Hawley, J. F. 2006, *ApJ*
- Schnittman, J. D., & Krolik, J. H. 2009, *ApJ* 701, 1175
- Schnittman, J. D., & Krolik, J. H. 2010, *ApJ*, 712, 908
- Schnittman, J. D., & Krolik, J. H. 2012, in preparation
- Shakura, N. I., & Sunyaev, R. A. 1973, *A&A*, 24, 337
- Shimura, T., & Takahara, F. 1995, *ApJ*, 445, 780
- Socrates, A., Davis, S. W., & Blaes, O. 2004, *ApJ*, 601, 405
- Socrates, A. 2010, *ApJ*, 719, 784
- Steiner, J. F., McClintock, J. E., Remillard, R. A., Gou, L., Yamada, S., & Narayan, R. 2010, *ApJ*, 718, L117
- Stern, B. E., Poutanen, J., Svensson, R., Sikora, M., & Begelman, M. C. 1995, *ApJ*, 449, L13
- Svensson, R., & Zdziarski, A. A. 1994, *ApJ*, 436, 599
- Tanaka, Y. et al. 1995, *Nature* 375, 659
- Titarchuk, L., & Shrader, C. 2002, *ApJ*, 567, 1057
- Turolla, R., Zane, S., & Titarchuk, L. 2002, *ApJ*, 576, 349
- Uttley, P., Wilkinson, T., Cassatella, P., Wilms, J., Pottschmidt, K., Hanke, M., & Bock, M. 2011, *MNRAS*, 414, L60
- Walton, D. J., Reis, R. C., Cackett, E. M., Fabian, A. C., & Miller, J. M. 2012, *MNRAS*, 422, 2510
- Wilkins, D. R., & Fabian, A. C. 2011, *MNRAS*, 414, 1269
- Zdziarski, A. A., Johnson, W. N., & Magdziarz, P. 1996, *MNRAS*, 283, 193
- Zhu, Y., Davis, S. W., Narayan, R., Kulkarni, A. K., Penna, R. F., & McClintock, J. E. 2012, *MNRAS* accepted, [arXiv:1202.1530]

